CLAIMS

We claim:

1. A speech processing board comprising:

multiple processor modules, each said processor module having an associated local memory, each said processor module hosting at least one instance of a speech application task;

a storage system for storing speech task data, said speech task data comprising language models and finite state grammars;

a local communications bus communicatively linking each said processor module through which each said processor module can exchange speech task data with said storage system; and,

a communications bridge to a host system, said communications bridge providing an interface to said local communications bus through which data can be exchanged between said processor modules and said host system.

2. The speech processing board of claim 1, wherein each said processor module comprises:

a central processing unit (CPU) core having at least one memory cache which can be accessed by said CPU core;

a processor bridge communicatively linking said CPU core to said local communications bus; and,

a memory controller through which said CPU core can access said local memory, said memory controller linked to said CPU core through a processor local bus.

3. The speech processing board of claim 2, further comprising a language model cache disposed in said local memory.

4. The speech processing board of claim 2, further comprising a finite state grammar table disposed in said local memory.

5. The speech processing board of claim 1, wherein said storage system comprises:

a fixed storage device accessible by said processor modules through said communications bridge, wherein said fixed storage device stores active language models and finite state grammars used by said speech application tasks hosted by said processor modules;

a commonly addressed language model cache, said language model cache storing at least one image of a language model stored in said fixed storage device, each said processor module accessing said language model cache through said communications bridge at a common address; and,

a boot memory storing initialization code, said boot memory communicatively linked to said processor modules through said communications bridge, each said processor module accessing said boot memory during an initial power-on sequence.

6. The speech processing board of claim 1, wherein said local communications bus is a PCI bus.

7. The speech processing board of claim 6, wherein said PCI bus is a 64-bit, 133MHz PCI bus.

8. The speech processing board of claim 6, wherein said PCI bus is a 64-bit, 66MHz PCI bus.

9. The speech processing board of claim 1, wherein said communications bridge comprises a PCI-to-PCI bridge having a PCI interface to said host system and an interface to an H.1x0 bus.

10. The speech processing board of claim 9, wherein said communications bridge further comprises a processing element for managing message communications between the speech processing board and said host system according to a messaging protocol provided by said host system.

11. The speech processing board of claim 1, wherein said communications bridge is implemented in a field programmable gate array (FPGA).

12. The speech processing board of claim 1, further comprising a serial audio channel communicatively linking said processor modules to said communications bridge, said serial audio channel providing a medium upon which audio data can be exchanged between individual processor modules and said communications bridge.

13. The speech processing board of claim 12, further comprising an audio stream processor coupled to said communications bridge, said audio stream processor configured to extract audio information received in said communications bridge, store said extracted audio information and distribute said audio information over said serial audio channel to selected ones of said processor modules based on hosted instances of speech applications in each said processor module.

14. The speech processing board of claim 12, further comprising an ethernet switch coupled to said communications bridge, said ethernet switch configured to transmit and receive packetized audio information to and from an external network.

15. The speech processing board of claim 1, wherein said host system is a CT media services system.

16. The speech processing board of claim 1, wherein said host system is a voice over IP (VoIP) gateway/endpoint.

17. A speech processing board comprising:

multiple processor modules in the speech processing board;

a PCI-to-PCI bridge interfacing said local PCI interface to a host CT system, said bridge comprising interfaces to an H.1x0 bus and a PCI bus;

a local PCI interface linking each said processor module to said PCI-to-PCI bridge;

a fixed storage communicatively linked to said PCI-to-PCI bridge and accessible by said processor modules through a drive controller;

a language model cache communicatively linked to said bridge; and,

a boot memory communicatively linked to said bridge, said boot memory storing initialization code.

18. A high-volume speech processing method comprising the steps of:

loading and executing a plurality of speech application tasks in selected ones of multiple processor modules in a speech processing board;

loading in a commonly addressed storage separate from said multiple processor modules selected language models for use by said speech application tasks;

receiving audio data over an audio channel and distributing said audio data to particular ones of said processor modules, wherein said distribution of said audio data to particular ones of said processor modules is determined based upon a speech application tasks executing in said particular ones of said processor modules;

processing said received audio data in said particular ones of said processor modules using said language models selected for use by said speech application tasks; and,

caching in said selected ones of said multiple processor modules portions of said selected language models used by said speech application tasks.

19. The speech processing method of claim 18, further comprising the steps of:

collecting speech task results from said selected ones of said multiple processor modules; and,

forwarding said collected speech task results to a host computer telephony (CT) system over a host communications bus.

20. A machine readable storage having stored thereon a computer program for processing speech, said computer program having a plurality of code sections executable by a machine for causing the machine to perform the steps of:

loading and executing a plurality of speech application tasks in selected ones of multiple processor modules in a speech processing board;

loading in a commonly addressed storage separate from said multiple processor modules selected language models for use by said speech application tasks;

receiving audio data over an audio channel and distributing said audio data to particular ones of said processor modules, wherein said distribution of said audio data

to particular ones of said processor modules is determined based upon a speech application tasks executing in said particular ones of said processor modules;

processing said received audio data in said particular ones of said processor modules using said language models selected for use by said speech application tasks; and,

caching in said selected ones of said multiple processor modules portions of said selected language models used by said speech application tasks.

21. The machine readable storage of claim 20, further comprising the steps of:

collecting speech task results from said selected ones of said multiple processor modules; and,

forwarding said collected speech task results to a host computer telephony (CT) system over a host communications bus.